

## **CLAIMS**

1. **(Withdrawn)** A method comprising:
  - receiving a request to be transmitted over a network to a particular client device;
  - determining an available downstream bandwidth that is associated with the client device;
  - determining an available upstream bandwidth that is associated with the client device, the available upstream bandwidth being different than the available downstream bandwidth;
  - determining a request data size associated with the request;
  - determining an estimated response data size associated with the request; and
  - in an event that the available downstream bandwidth is greater than the request data size and the available upstream bandwidth is greater than the estimated response data size, transmitting the request to the client device.
2. **(Withdrawn)** The method as recited in claim 1 wherein the client device comprises a television set-top box.
3. **(Withdrawn)** The method as recited in claim 1 wherein the network comprises a cable television network.

4. **(Withdrawn)** The method as recited in claim 1 wherein the available upstream bandwidth is less than the available downstream bandwidth.

5. **(Withdrawn)** The method as recited in claim 1 wherein the available upstream bandwidth is configurable.

6. **(Withdrawn)** The method as recited in claim 1 wherein the available upstream bandwidth is configurable to vary over time.

7. **(Withdrawn)** The method as recited in claim 1 wherein the available downstream bandwidth is configurable.

8. **(Withdrawn)** The method as recited in claim 1 wherein the available downstream bandwidth is configurable to vary over time.

9. **(Withdrawn)** A cable television system headend configured to perform the method as recited in claim 1.

10. **(Withdrawn)** One or more computer-readable media comprising computer-executable instructions that direct a computing system to perform the method as recited in claim 1.

11. **(Withdrawn)** A method comprising:

receiving a request to be transmitted over a downstream network path to a particular client device, the client device configured to transmit a response over an upstream network path;

identifying one or more other client devices that contend for the upstream network path; and

in an event that there are no outstanding requests to any of the one or more other client devices, transmitting the request to the particular client device.

12. **(Withdrawn)** The method as recited in claim 11 wherein an available bandwidth on the upstream network path differs from an available bandwidth on the downstream network path.

13. **(Withdrawn)** The method as recited in claim 11 wherein the client device comprises a television set-top box.

14. **(Withdrawn)** The method as recited in claim 11 wherein the downstream network path is part of a cable television network.

15. **(Withdrawn)** The method as recited in claim 11 further comprising:

determining a maximum number of contending client devices associated with the upstream network path to which outstanding requests may be pending;

receiving a second request to be transmitted over the downstream network path to a second client device, the second client device configured to transmit a response to the second request over the upstream network path;

determining a number of other client devices associated with the upstream network path to which pending requests are outstanding; and

in an event that the number of other client devices to which pending requests are outstanding is less than the maximum number of contending client devices to which outstanding requests may be pending, transmitting the second request to the second client device.

16. **(Withdrawn)** One or more computer-readable media comprising computer-executable instructions that direct a computing system to perform the method as recited in claim 11.

17. **(Withdrawn)** A method comprising:

receiving a request to be transmitted over a network to a particular client device;

and

placing the request in a request queue, the request queue being associated with an upstream network channel over which a response to the request may be received from the client device.

18. **(Withdrawn)** The method as recited in claim 17 further comprising:

transferring the request from the request queue to a dispatch queue when an available downstream network bandwidth is sufficient to accommodate the request and an available upstream network bandwidth is sufficient to accommodate an anticipated response to the request.

19. **(Withdrawn)** One or more computer-readable media comprising computer-executable instructions that direct a computing system to perform the method as recited in claim 17.

20. (Currently Amended) A method comprising:

identifying a plurality of upstream network paths associated with a single downstream network path;

establishing a plurality of application-upstream path request queues, wherein first and second application-upstream path request queues of the plurality of application upstream path request queues are associated with first and second upstream network paths of the plurality of upstream network paths, respectively;

receiving from one or more applications, first and second requests to be transmitted to first and second client devices, respectively;

determining that the first client device is associated with the first upstream network path and that the second client device is associated with the second upstream network path;

storing the first request in the first application-upstream path request queue and storing the second request in the second application-upstream path request queue; and

performing a scheduling process, the scheduling process comprising:

verifying that a bandwidth available on the downstream network path is larger than a data size of the first request;

verifying that a bandwidth available on the first upstream network path is larger than a data size of an anticipated response to the first request; and

moving the first request to a first dispatch queue.

21. **(Currently Amended)** The method as recited in claim 20 wherein the scheduling process further comprises:

verifying that the bandwidth available on the downstream network path is larger than a combined data size of the first request and the second request combined;

verifying that a bandwidth available on the second upstream network path is larger than a data size of an anticipated response to the second request; and

moving the second request to a second dispatch queue.

22. **(Original)** The method as recited in claim 21 further comprising:

transmitting the first and second requests from the first and second dispatch queues over the downstream network path.

23. **(Withdrawn)** A method comprising:

establishing an application request channel;

assigning to the application request channel a portion of an available downstream bandwidth; and

associating an application with the application request channel such that a request initiated by the application can only be transmitted when the portion of the available downstream bandwidth is large enough to accommodate the request.

24. **(Withdrawn)** The method as recited in claim 23 wherein the portion of the available downstream bandwidth comprises 100 percent of the available downstream bandwidth.

25. **(Withdrawn)** The method as recited in claim 23 further comprising:  
assigning to the application request channel a portion of an available upstream bandwidth such that the request initiated by the application can only be transmitted when the portion of the available upstream bandwidth is large enough to accommodate an anticipated response to the request.

26. **(Withdrawn)** The method as recited in claim 23 wherein the portion of the available upstream bandwidth comprises 100 percent of the available upstream bandwidth.

27. **(Withdrawn)** One or more computer-readable media comprising computer-executable instructions that direct a computing system to perform the method as recited in claim 23.

28. **(Withdrawn)** A method comprising:  
establishing an application request channel;  
assigning to the application request channel a portion of an available upstream bandwidth; and  
associating an application with the application request channel such that a request initiated by the application can only be transmitted when the portion of the available upstream bandwidth is large enough to accommodate an anticipated response to the request.

29. **(Withdrawn)** The method as recited in claim 28 wherein the portion of the available upstream bandwidth comprises 100 percent of the available upstream bandwidth.

30. **(Withdrawn)** The method as recited in claim 28 further comprising: assigning to the application request channel a portion of an available downstream bandwidth such that the request initiated by the application can only be transmitted when the portion of the available downstream bandwidth is large enough to accommodate the request.

31. **(Withdrawn)** The method as recited in claim 28 wherein the portion of the available downstream bandwidth comprises 100 percent of the available downstream bandwidth.

32. **(Withdrawn)** One or more computer-readable media comprising computer-executable instructions that direct a computing system to perform the method as recited in claim 28.

33. **(Currently Amended)** A system comprising:

an application request queue for receiving requests to be transmitted to client devices via a downstream network path;

a plurality of upstream request queues for grouping the requests based on upstream network paths over which responses to the requests may be received; and

a bandwidth management system configured to schedule the requests from the plurality of upstream request queues for transmission to the client devices based, at least in part, on available bandwidth associated with one or more of the upstream network paths.

34. **(Original)** The system as recited in claim 33 wherein the bandwidth management system is further configured to schedule the requests from the plurality of upstream request queues for transmission to the client devices based on available bandwidth associated with the downstream network path.

35. **(Original)** The system as recited in claim 33, implemented as a component of a cable television system headend.

36. **(Currently Amended)** A system comprising:

one or more application request channels for receiving requests to be transmitted to client devices via a downstream network path, individual ones of the application request channels having an associated downstream path bandwidth allocation;

a plurality of upstream request queues for grouping the requests based on upstream network paths over which responses to the requests may be received; and

a bandwidth management system configured to schedule the requests from the plurality of upstream request queues for transmission to the client devices based, at least in part, on a request size and the downstream path bandwidth allocation associated with the application request channel through which the request was received.

37. **(Currently Amended)** The system as recited in claim 36 wherein individual ones of the application request channels have an associated upstream path bandwidth allocation and wherein the bandwidth management system is further configured to schedule the requests from the plurality of upstream request queues based on an estimated response size and the upstream path bandwidth allocation associated with the application request channel through which the request was received.

38. **(Original)** The system as recited in claim 36, implemented as a component of a cable television system headend.

39. **(Currently Amended)** A system comprising:

one or more application request channels for receiving requests to be transmitted to client devices via a downstream network path, individual ones of the application request channels having an associated upstream path bandwidth allocation;

a plurality of upstream request queues for grouping the requests based on upstream network paths over which responses to the requests may be received; and

a bandwidth management system configured to schedule the requests from the plurality of upstream request queues for transmission to the client devices based, at least in part, on an anticipated request response size and the upstream path bandwidth allocation associated with the application request channel through which the request was received.

40. **(Currently Amended)** The system as recited in claim 39 wherein individual ones of the application request channels have an associated downstream path bandwidth allocation and wherein the bandwidth management system is further configured to schedule the requests from the plurality of upstream request queues based on a request size and the downstream path bandwidth allocation associated with the application request channel through which the request was received.

41. **(Original)** The system as recited in claim 39, implemented as a component of a cable television system headend.

42. **(Currently Amended)** A system, comprising:

means for receiving a request from an application, the request to be transmitted to a particular client device;

means for queuing the request for transmission over a downstream network path based on identification of an upstream network path over which a response to the request may be received; and

means for scheduling transmission of the request based, ~~at least in part~~, on an available upstream network bandwidth.

43. **(Currently Amended)** The system as recited in claim 42 wherein ~~the~~ means for scheduling further comprises means for scheduling transmission of the request is further based, ~~at least in part~~, on an available downstream network bandwidth.

44. **(Currently Amended)** The system as recited in claim 42 wherein ~~the~~ means for scheduling further comprises means for scheduling transmission of the request is further based, ~~at least in part~~, on a number of pending requests previously transmitted to other client devices associated with the upstream network path.

45. **(Original)** The system as recited in claim 42 further comprising means for estimating an anticipated response size.

46. **(Currently Amended)** One or more computer-readable storage media comprising computer-executable instructions that, when executed, direct a computing system to perform a method comprising:

receiving a request;  
identifying an upstream network path associated with the request;  
queuing the request for transmission over a downstream network path based on the upstream network path that was identified; and  
scheduling the request for transmission based on an available bandwidth associated with the upstream network path.

47. **(Currently Amended)** One or more computer-readable storage media comprising computer-executable instructions that, when executed, direct a computing system to perform a method comprising:

receiving a request;  
identifying a client device to which the request is to be transmitted;  
identifying an upstream network path associated with the client device;  
identifying one or more other client devices that are associated with the upstream network path;  
queuing the request for transmission over a downstream network path based on the upstream network path that was identified; and  
scheduling the request for transmission based on a number of the one or more other client devices to which pending requests have been transmitted.

48. **(Currently Amended)** One or more computer-readable storage media comprising computer-executable instructions that, when executed, direct a computing system to perform a method comprising:

receiving a request through an application request channel;

identifying a client device to which the request is to be transmitted;

identifying an upstream network path associated with the client device;

queuing the request for transmission over a downstream network path based on the upstream network path that was identified;

identifying an upstream bandwidth allocation associated with the application request channel; and

scheduling the request for transmission based on a comparison between a portion of an available bandwidth associated with the upstream network path, the upstream bandwidth allocation, and a size of an anticipated response to the request, wherein the portion of the available bandwidth is determined based on the upstream bandwidth allocation associated with the application request channel.

49. **(Currently Amended)** The one or more computer-readable storage media as recited in claim 48, wherein the method further comprises estimating [[a]] the size of [[an]] the anticipated response to the request.

50. **(Currently Amended)** One or more computer-readable storage media comprising computer-executable instructions that, when executed, direct a computing system to perform a method comprising:

receiving a request through an application request channel;

identifying a client device to which the request is to be transmitted;

identifying an upstream network path associated with the client device;

identifying a downstream network path associated with the client device;

queuing the request for transmission over the downstream network path~~based on the upstream network path~~; and

identifying a downstream bandwidth allocation associated with the application request channel;

~~scheduling the request for transmission based on a comparison between a portion of an available bandwidth associated with the downstream network path, the downstream bandwidth allocation, and a size of the request, wherein the portion of the available bandwidth associated with the downstream network path is determined based on the downstream bandwidth allocation associated with the application request channel.~~

51. **(Currently Amended)** One or more computer-readable storage media comprising computer-executable instructions that, when executed, direct a computing system to perform a method comprising:

scheduling requests for transmission, wherein the scheduling is performed at a scheduling interval by placing the requests in a dispatch queue, the scheduling interval defined by a particular value;

determining a number of requests in the dispatch queue; and

in an event that the number of requests in the dispatch queue is greater than or equal to a maximum number of allowed pending requests, preventing additional requests from being scheduled for transmission until the number of requests in the dispatch queue is less than the maximum number of allowed pending requests.

52. **(Previously Presented)** The one or more computer-readable storage media as recited in claim 51, wherein the particular value is configurable.

53. **(Previously Presented)** The one or more computer-readable storage media as recited in claim 51, wherein the maximum number of allowed pending requests is equal to an integer multiple of a number of requests that can be scheduled into the dispatch queue during a single scheduling interval.